Evaluation and optimization of the I/O scalability of the (Open)IFS atmospheric model using XIOS

Xavier Yepes-Arbós
Mario C. Acosta

The research leading to these results has received funding from the EU H2020 Framework Programme under grant agreement no. 823988
This material reflects only the author's view and the Commission is not responsible for any use that may be made of the information it contains

8/9/2020

11th JLESC Workshop

Introduction

- Exascale supercomputers will allow Earth System Models (ESMs) to make simulations at an unprecedented level of horizontal resolution.
- But this has implications:
  - A huge amount of data will be generated that must be efficiently written into the storage system.
  - No more offline post-processing is affordable due to the size of the "raw" data.
  - A high cost of storage systems due to the huge data size.

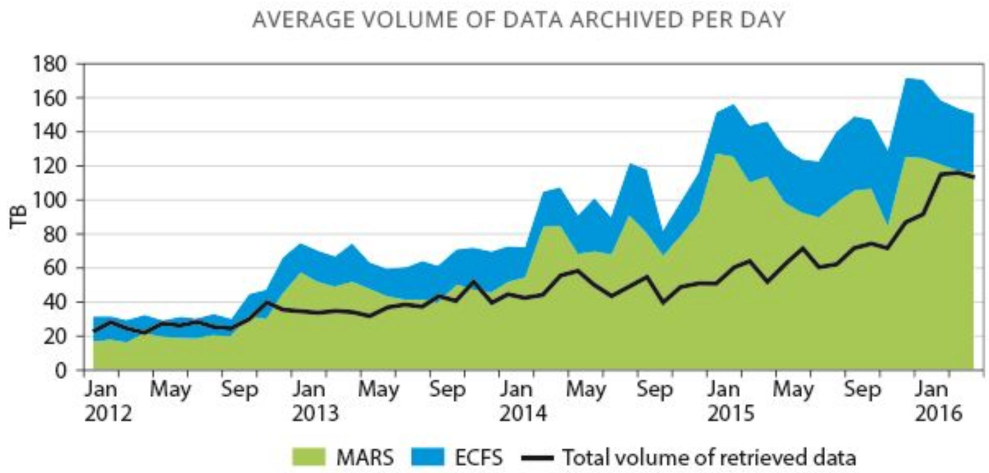

Introduction

- Not much attention was paid on improving I/O of ESMs because it did not use to be an issue.
- This was the case of Numerical Weather Prediction (NWP) models such as IFS, where one of its output schemes uses sequential I/O.
- Sequential I/O is not scalable for such high resolution grids, and even less, for future exascale machines.

(Open)IFS overview

- The Integrated Forecasting System (IFS) is a global data assimilation and forecasting system which includes the modelling of the atmospheric composition developed by the European Centre for Medium-Range Weather Forecasts (ECMWF).
- It has two different output schemes:
  - The Météo-France (MF) I/O server which is fast and efficient from a computational point of view. It is only used at ECMWF, such its operational forecasts.
  - A sequential I/O scheme which is slow and inefficient from a computational point of view. It is used by non-ECMWF users, this is, in OpenIFS.
- OpenIFS is a free and simplified version of IFS available under a license.

Objective: Integrate XIOS

- The I/O issue is typically addressed by adopting scalable parallel I/O solutions.

- In the climate community, a widely I/O tool used is XIOS.

- The XML Input/Output Server (XIOS) is an asynchronous MPI parallel I/O server developed by the Institute Pierre Simon Laplace (IPSL).

- XIOS has the following features needed for climate modelling:
  - Output files are in netCDF format.
  - Written data is CMIP-compliant (CMORized).
  - It is able to post-process data online to generate diagnostics.

XIOS: Some technical features

- From a computational point of view, XIOS is thought to address:
  - The inefficient legacy read/write process.
  - The unmanageable size of “raw” data.

- By implementing:
  - Scalable parallel I/O.
  - Online post-processing.

- But it has been only tested for petascale supercomputers, so it is necessary to:
  - Stress different aspects such as memory consumption, MPI scalability, netCDF parallel I/O or **data compression**.

(Open)IFS-XIOS integration scheme

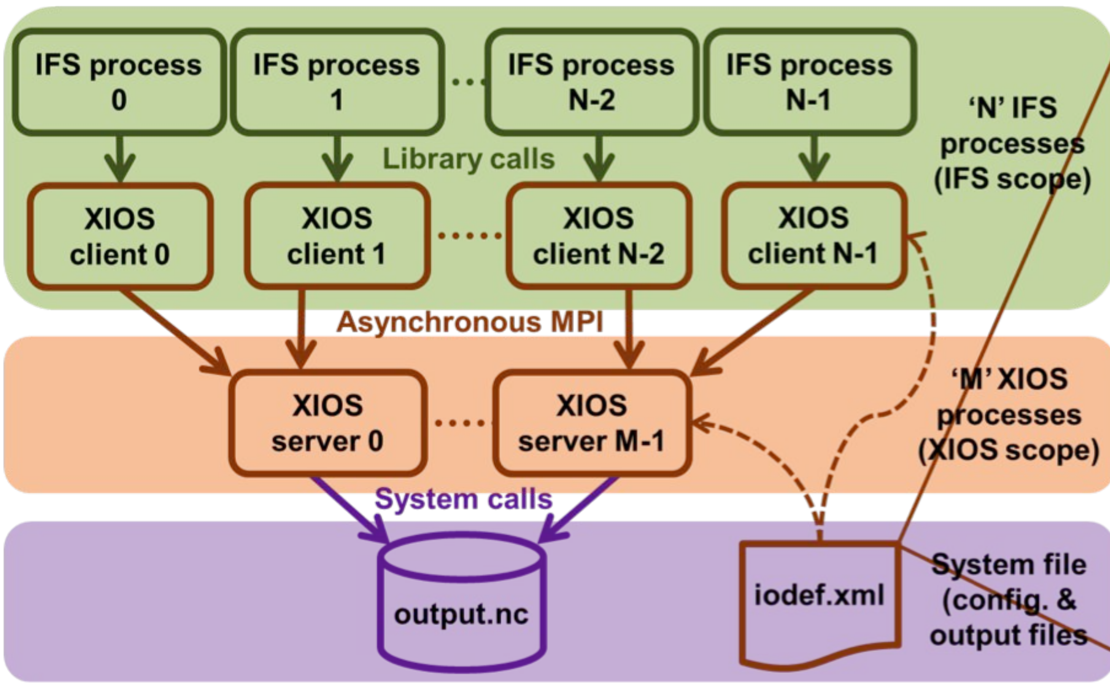


```
<grid id="grid_3d">
  <domain id="domain_2d"
  type="rectilinear" ni_glo=... />
  <axis id="axis_1d" n_glo=... />
</grid>

<field_definition>
  <field id="temp" grid_ref="grid_3d" />
  <field id="pressure" grid_ref="grid_3d" />
</field_definition>

<file_definition>
  <file name="output" freq_output="1mo">
    <field field_ref="temp" />
    <field field_ref="pressure" />
  </file>
</file_definition>
```

(Open)IFS-XIOS integration summary

- Scientific highlights:
  - Both grid-point and spectral fields are supported.
  - All surface and 3D fields can be output.
  - Different vertical levels are available: model, pressure, theta and PV levels.
  - No longer needed to set up the FullPos namelist (NAMFPC).
  - FullPos spectral fitting is available.
  - Physical tendencies and fluxes output (PEXTRA fields) are also supported.

- Both XIOS 2.0 and 2.5 versions have been tested.

(Open)IFS-XIOS integration summary

- Highlights from the computational performance point of view:
  - In-depth benchmarking: the **overhead** of outputting data through XIOS is really **small** if using enough computational resources.
  - A profiling and performance analysis was done to detect potential bottlenecks.
  - Two different optimizations are available (switchable in the XIOS XML namelist):
    - Computation and communication overlap.
    - Sends from (Open)IFS to XIOS either in double or single precision.

- Different XIOS features available (listed only some of them):
  - Horizontal interpolations (from reduced Gaussian to rectangular Gaussian).
  - Arithmetic operations.
  - Time operations: average, maximum, minimum, etc.
  - **Lossless data compression** using gzip through HDF5.

Computational performance of IFS-XIOS

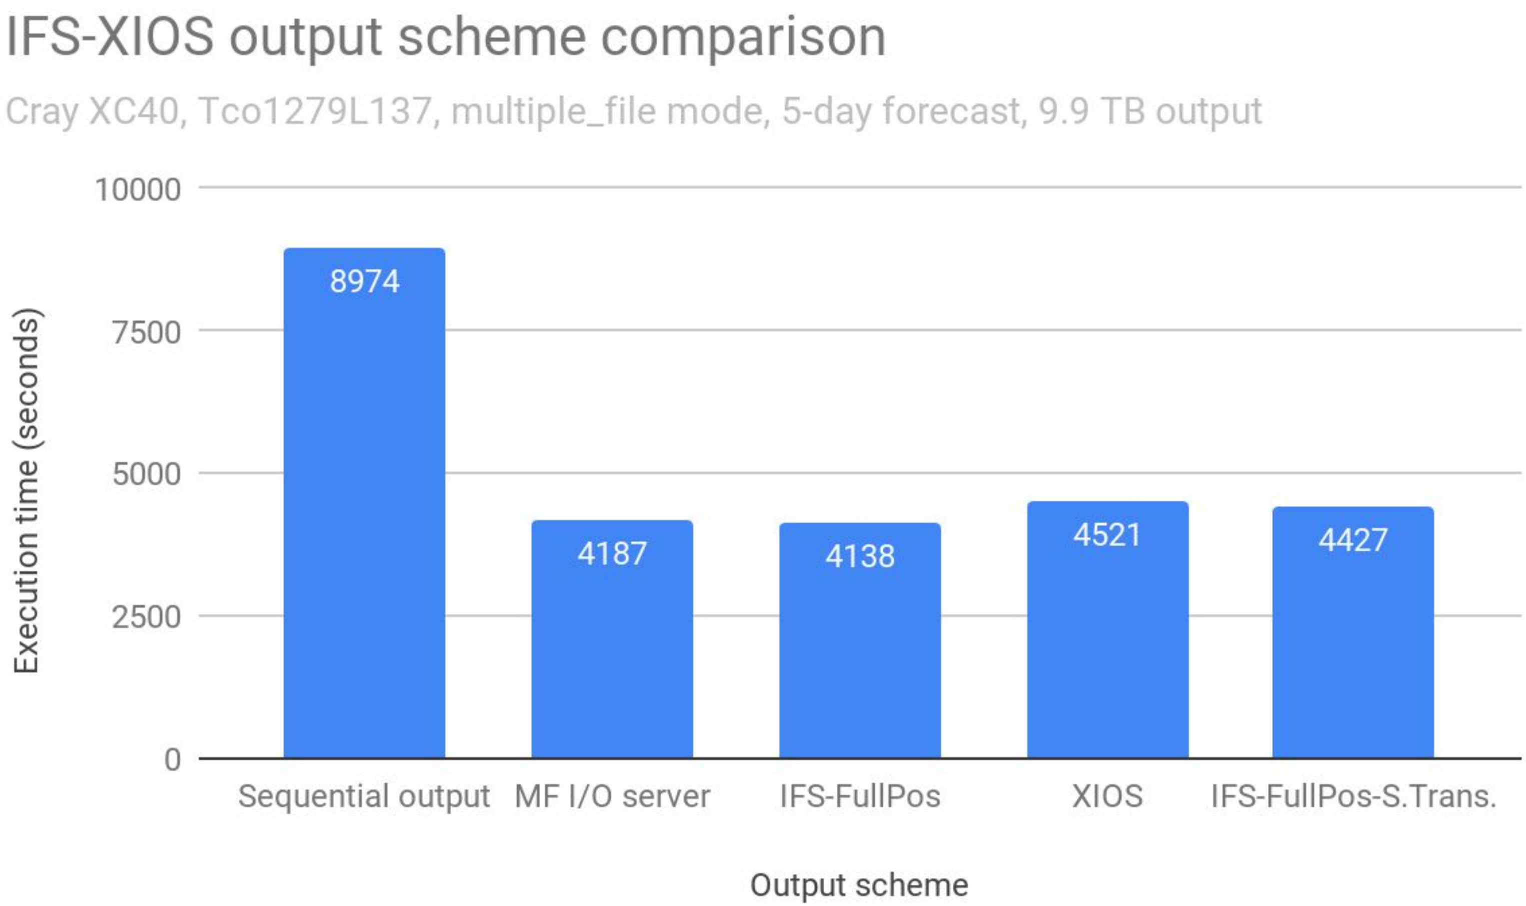

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

XIOS performance issue

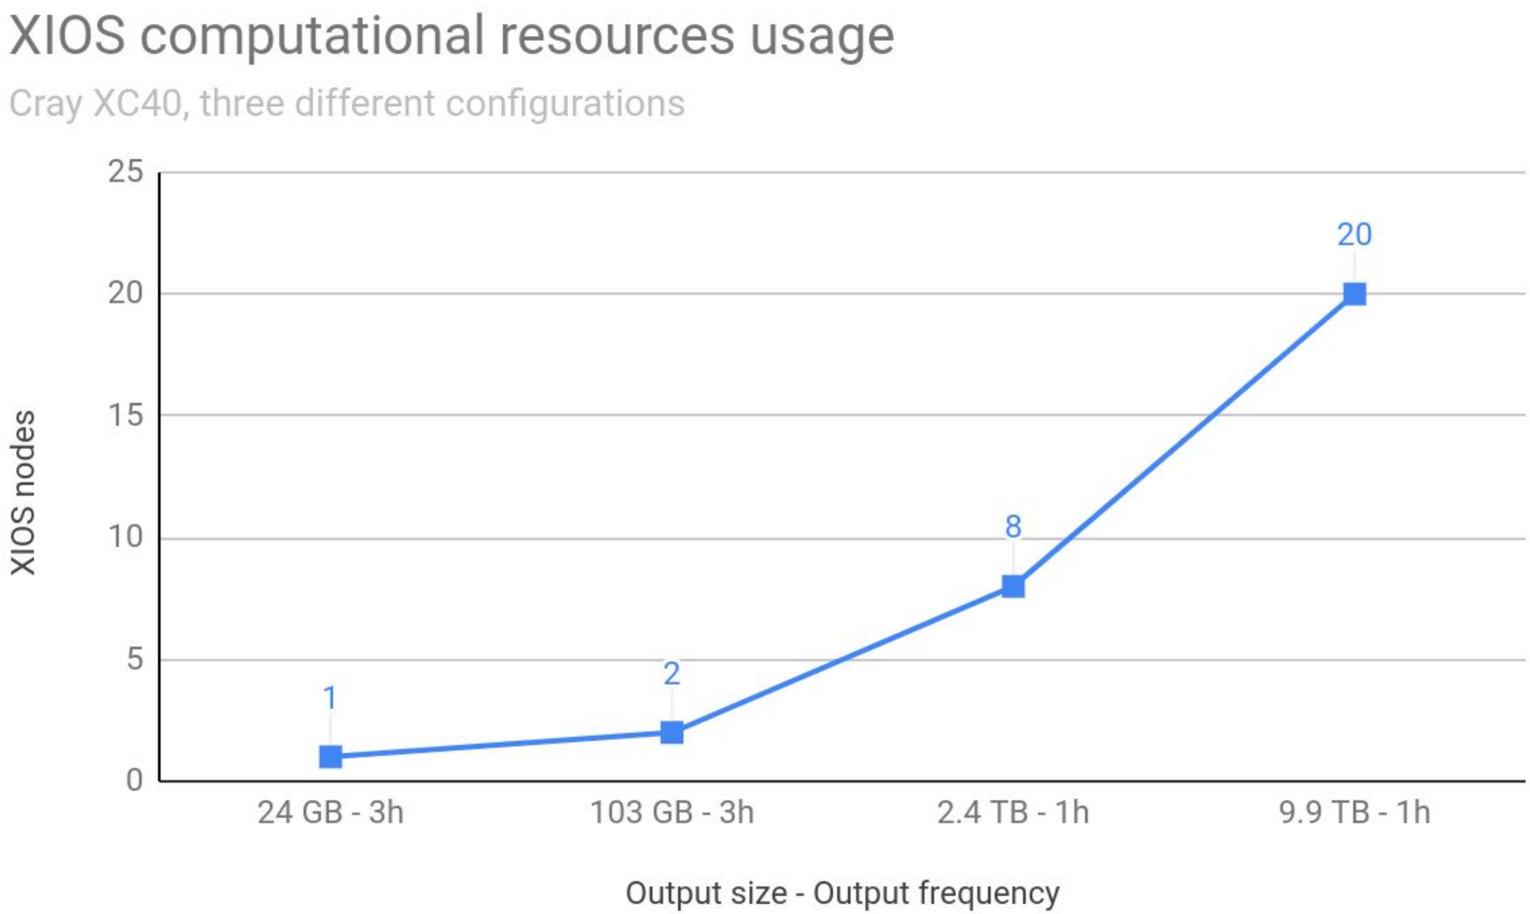

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

XIOS performance issue

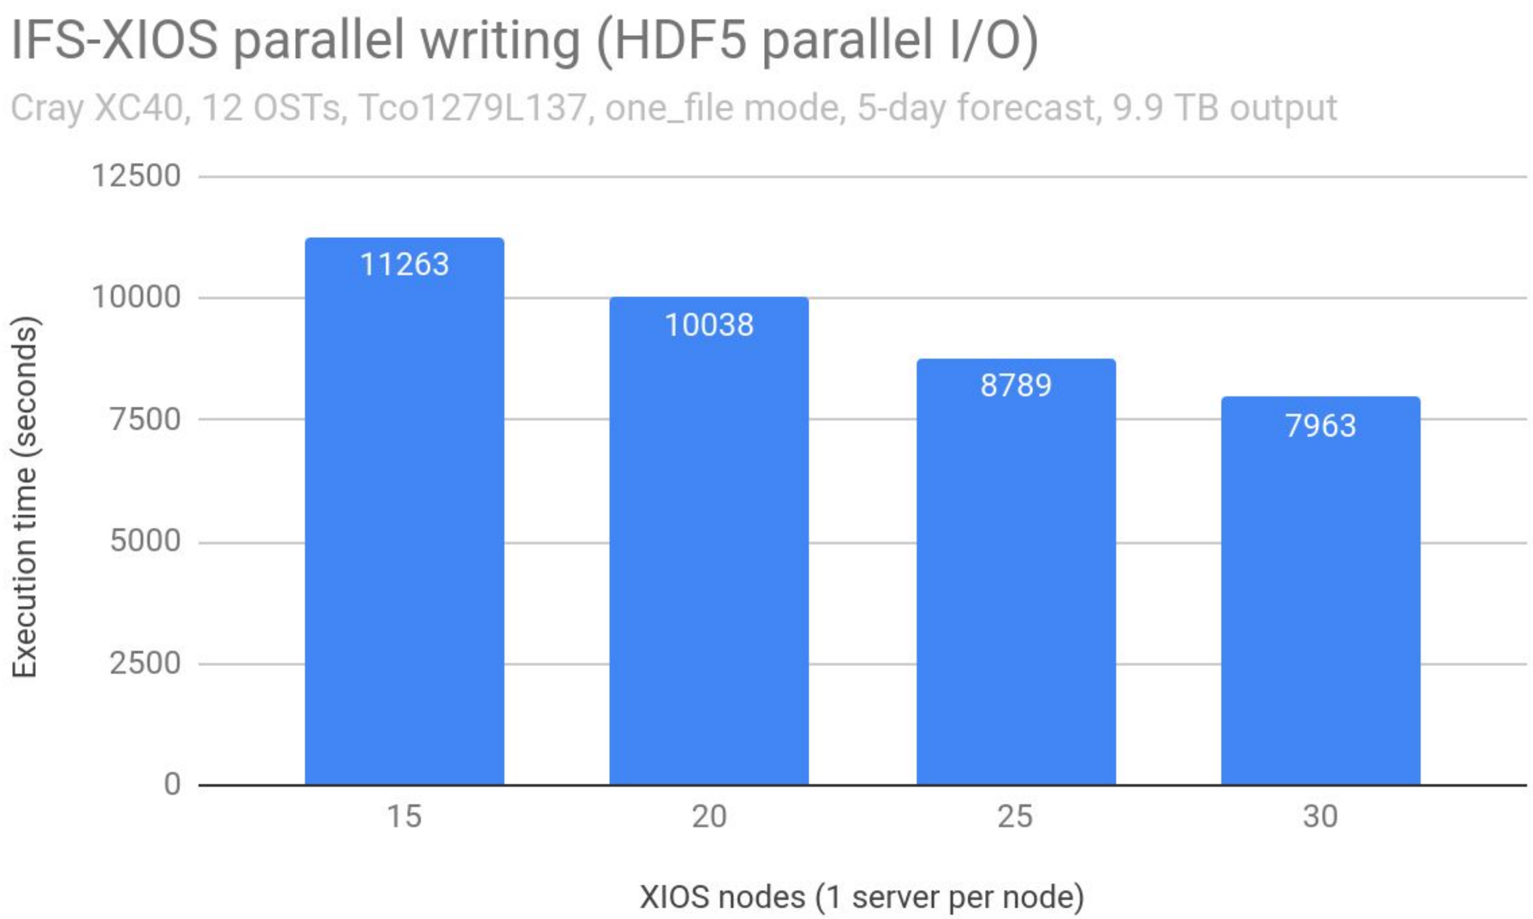

What about XIOS compression?

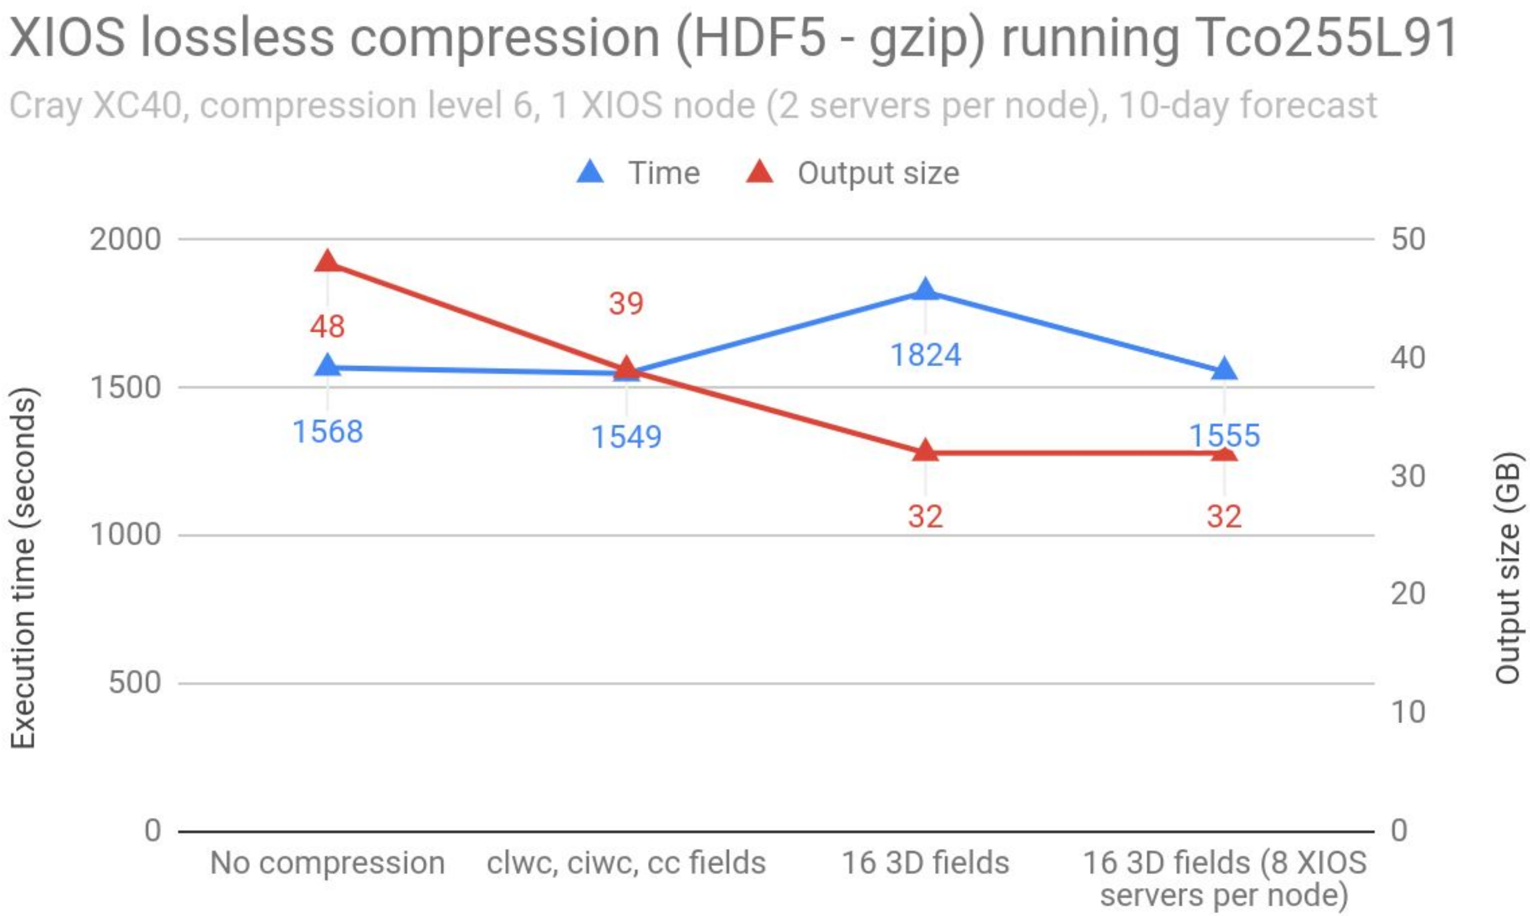

What about XIOS compression?

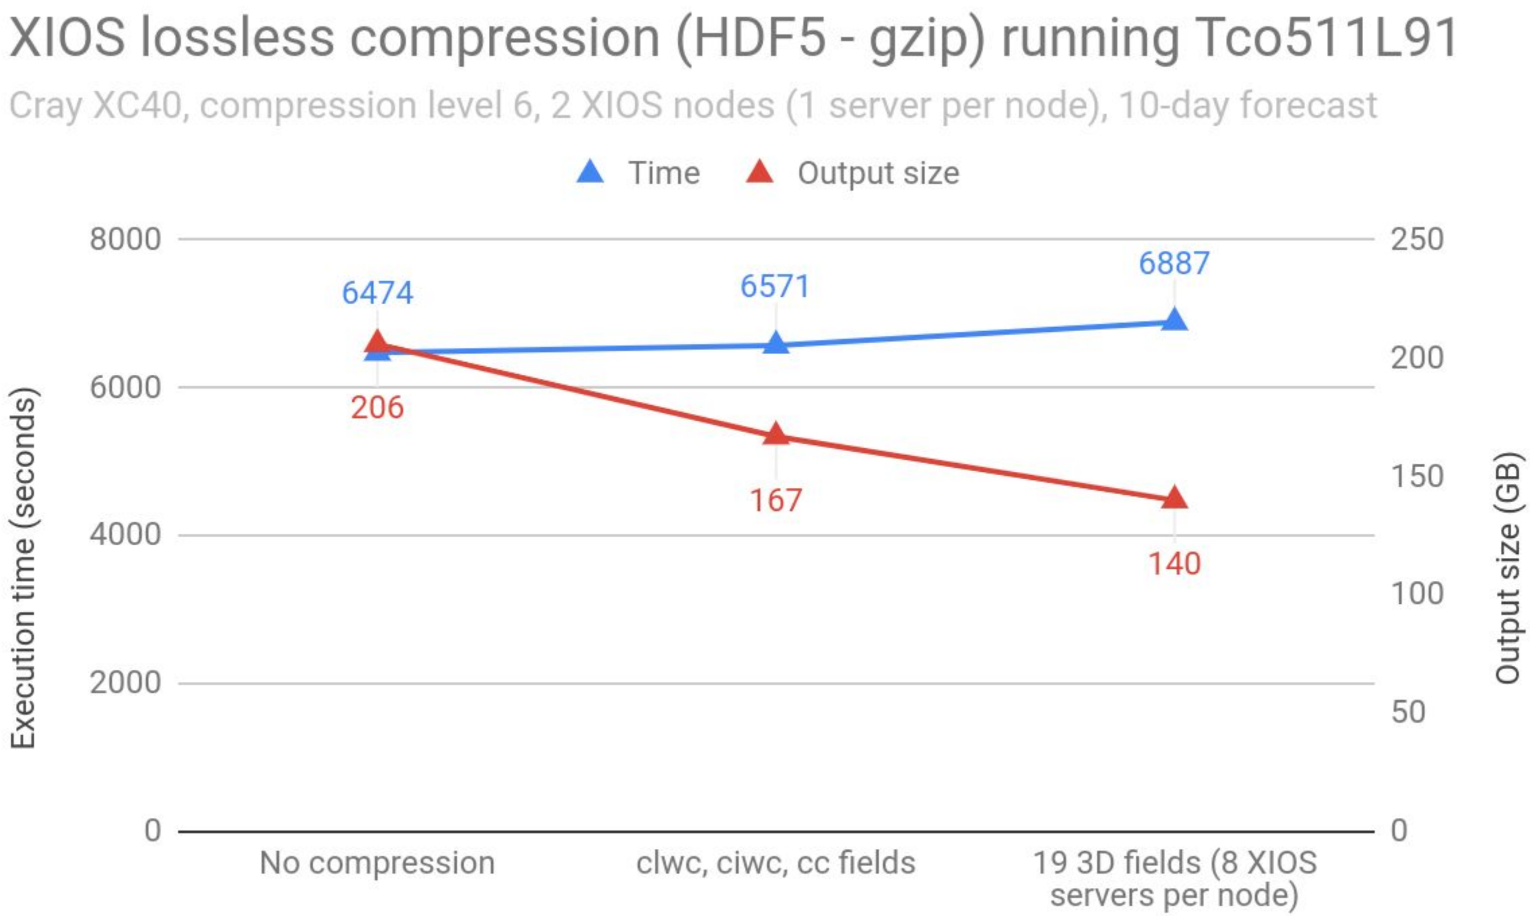

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

What about XIOS compression?

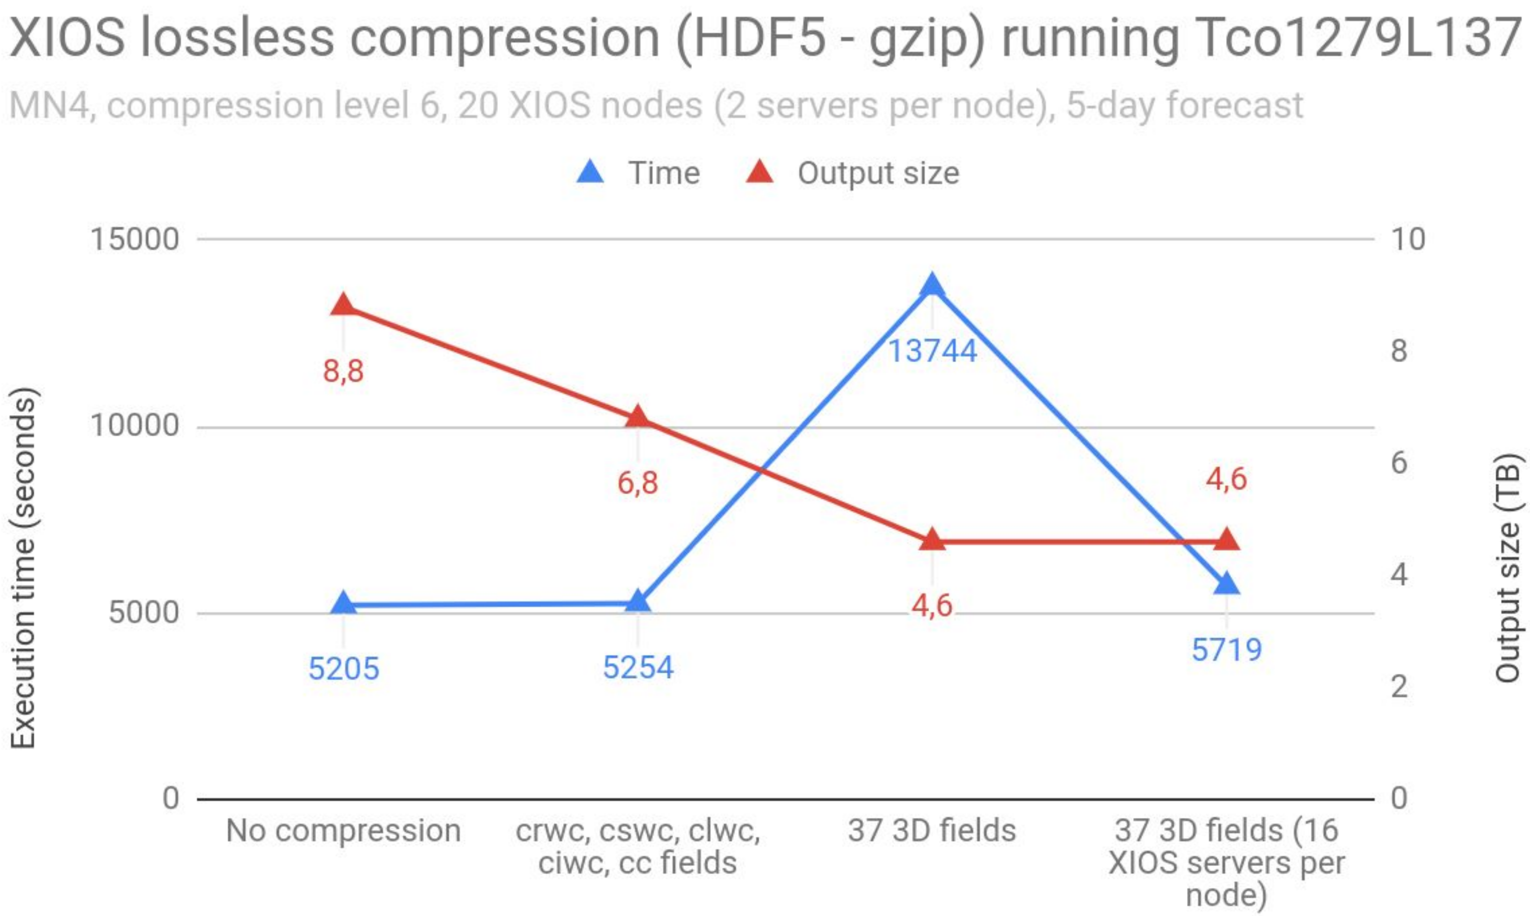

Lossy compression filter for XIOS?

- The default lossless compression filter of HDF5 does not fit our needs:
  - If compression ratio is high, it takes too much time.
  - If it takes a reasonable amount of time, compression ratio is not enough.

- We want to explore if lossy compression is adequate for climate modelling. In particular, we would consider to use lossy compression in XIOS if it fulfills the following points:
  - Reach high compression ratios.
  - Enough compression speed to considerably mitigate the I/O overhead.
  - Keep high accuracy.

- In addition, it would be interesting to have OpenMP support as well as be integrated with HDF5 parallel I/O.

Open questions and collaboration opportunities

- What types of lossy compression are more suitable for climate modelling? Can we use same compression for all variables?

- In particular, we are interested in the SZ compressor from ANL. Do you think it is suitable for our needs?

- Thus, might there be a potential collaboration between ANL and BSC?

- The SZ compressor is already registered as a third-party filter of HDF5. We would like to explore with ANL if it is enough for XIOS, or we would need to develop a particular solution. Is it compatible with experimental HDF5 parallel I/O? And OpenMP?

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Thank you

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

xavier.yepes@bsc.es